

# Is This the Right Place?

## Geometric-Semantic Pose Verification for Indoor Visual Localization

Hajime Taira<sup>1</sup> Ignacio Rocco<sup>2</sup> Jiri Sedlar<sup>3</sup> Masatoshi Okutomi<sup>1</sup> Josef Sivic<sup>2,3</sup>  
 Tomas Pajdla<sup>3</sup> Torsten Sattler<sup>4</sup> Akihiko Torii<sup>1</sup>

<sup>1</sup>Tokyo Institute of Technology, <sup>2</sup>Inria, <sup>3</sup>CIIRC CTU in Prague, <sup>4</sup>Chalmers University of Technology

### Abstract

Visual localization in large and complex indoor scenes, dominated by weakly textured rooms and repeating geometric patterns, is a challenging problem with high practical relevance for applications such as Augmented Reality and robotics. To handle the ambiguities arising in this scenario, a common strategy is, first, to generate multiple estimates for the camera pose from which a given query image was taken. The pose with the largest geometric consistency with the query image, e.g., in the form of an inlier count, is then selected in a second stage. While a significant amount of research has concentrated on the first stage, there has been considerably less work on the second stage. In this paper, we thus focus on pose verification. We show that combining different modalities, namely appearance, geometry, and semantics, considerably boosts pose verification and consequently pose accuracy. We develop multiple hand-crafted as well as a trainable approach to join into the geometric-semantic verification and show significant improvements over state-of-the-art on a very challenging indoor dataset.

### 1. Introduction

Visual localization is the problem of estimating the 6 Degree-of-Freedom (DoF) pose from which an image was taken with respect to a 3D scene. Visual localization is vital to applications such as Augmented and Mixed Reality [15, 44], intelligent systems such as self-driving cars and other autonomous robots [41], and 3D reconstruction [61].

State-of-the-art approaches for accurate visual localization are based on matches between 2D image and 3D scene coordinates [10, 11, 16, 40, 46, 47, 56, 71, 72, 85]. These 2D-3D matches are either established using explicit feature

\*WILLOW project, Departement d'Informatique de l'École Normale Supérieure, ENS/INRIA/CNRS UMR 8548, PSL Research University.

<sup>†</sup>CIIRC - Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague.




Figure 1: **Using further modalities for indoor visual localization.** Given a set of camera pose estimates for a query image (a, g), we seek to identify the most accurate estimate. (b, h) Due to severe occlusion and weak textures, a state-of-the-art method [72] fails to identify the correct camera pose. To overcome those difficulties, we use several modalities along with visual appearance: (top) surface normals and (bottom) semantics. (c, i) Our approach verifies the estimated pose by comparing the semantics and surface normals extracted from the query (d, j) and database (f, l).

matching [40, 56, 71, 72, 85] or via learning-based scene coordinate regression [10, 11, 16, 46, 47, 65]. At large scale or in complex scenes with many repeating structural elements, establishing unique 2D-3D matches becomes a hard problem due to global ambiguities [40, 71, 85]. A strategy to avoid such ambiguities is to restrict the search space for 2D-3D matching. For example, image retrieval [49, 69] can

be used to identify a few parts of the scene most likely to be seen in a query image [29, 58, 72]. 2D-3D matching is then performed for each such retrieved place, resulting in one pose hypothesis per place. Subsequently, the “best” pose is selected as the final pose estimate for the query image.

Traditionally, the “best” pose has been defined as the pose with the largest number of inlier matches [29, 61]. Yet, it has been shown that a (weighted) inlier count is not a good decision criterion in the presence of repeated structures and global ambiguities [55]. Rather than only accounting for positive evidence in the form of the inliers, [72] proposed to compare the query photo against an image of the scene rendered using the estimated pose. [72] have shown that such a pixel-wise comparison, termed Dense Pose Verification (DensePV), leads to a significantly better definition of the “best” pose and subsequently improves pose accuracy.

In this paper, we follow the approach from [72], which focuses purely on comparing low-level appearance and geometry information between the re-rendered and actual query image. This paper asks the question whether it is possible to improve the pose verification stage and thus the pose accuracy of visual localization approaches. To this end, we analyze the impact of using further geometric and semantic modalities as well as learning in the verification stage. In detail, this paper makes the following contributions: **1)** We investigate the impact of using multiple modalities during pose verification for indoor visual localization in challenging scenarios. We hand-design several modifications of the original DensePV approach that integrate additional 3D geometry as well as normal and semantic information. We show that these approaches improve upon the original DensePV strategy, setting a new state-of-the-art in localization performance on the highly challenging InLoc dataset [72]. None of these approaches require fine-tuning on the actual dataset used for localization and are thus generally applicable. We are not aware of prior work that combines such modalities. **2)** We also investigate a trainable approach for pose verification. We show that it outperforms the original DensePV, which uses a hand-crafted representation. However, it is not able to outperform our novel modifications, even though it is trained on data depicting the scene used for localization. **3)** We show that there is still significant room for improvement by more advanced combinations, opening up avenues for future work. In addition, we show that a standard approach for semantic pose verification used for outdoor scenes in the literature [21, 22, 74, 75] is not applicable for indoor scenes. **4)** We make our source code and training data publicly available<sup>1</sup>.

## 2. Related Work

**Visual localization.** Structure-based visual localization uses a 3D scene model to establish 2D-3D matches be-

tween pixel positions in the query image and 3D points in the model [11, 14, 16, 29, 40, 42, 56, 65, 71, 72, 85]. The scene model can be represented either explicitly, *e.g.*, a Structure-from-Motion (SfM) point cloud [29, 39, 40, 56] or a laser scan [72], or implicitly, *e.g.*, through a convolutional neural network (CNN) [10, 11, 45] or a random forest [16, 45–47, 65]. In the former case, 2D-3D matches are typically established by matching local features such as SIFT [43]. In contrast, methods based on implicit scene representations directly regress 3D scene coordinates from 2D image patches [10, 11, 16, 45]. In both cases, the camera pose is estimated from the resulting 2D-3D matches by applying an  $n$ -point-pose solver [27, 35, 38] inside a RANSAC [18, 19, 25] loop. Methods based on scene coordinate regressions are significantly more accurate than approaches based on local features [11, 72]. Yet, they currently do not scale to larger and more complex scenes [72].

Closely related to visual localization is the *place recognition problem* [1, 4, 17, 26, 34, 36, 55, 60, 77, 78, 84]. Given a database of geo-tagged images, place recognition approaches aim to identify the place depicted in a given query image, *e.g.*, via image retrieval [3, 19, 49, 69, 76]. The geo-tag of the most similar database image is then often used to approximate the pose of the query image [30, 31, 77, 84]. Place recognition approaches can also be used as part of a visual localization pipeline [13, 29, 53, 62, 72]: 2D-3D matching can be restricted to the parts of the scene visible in a short list of  $n$  visually similar database images, resulting in one pose estimate per retrieved image. This restriction helps to avoid global ambiguities in a scene, *e.g.*, caused by similar structures found in unrelated parts of a scene, during matching [54]. Such retrieval-based methods currently constitute the state-of-the-art for large-scale localization in complex scenes [53, 55, 72]. In this paper, we follow this strategy. However, unlike previous work focused on improving the retrieval [1, 4, 20, 26, 34, 36] or matching [56, 72], we focus on the pose verification stage, *i.e.*, the problem of selecting the “best” pose from the  $n$  estimated poses.

An alternative to the localization approaches outlined above is to train a CNN that directly regresses the camera pose from a given input image [9, 12, 32, 33, 50, 79]. However, it was recently shown that such methods do not consistently outperform a simple image retrieval baseline [59].

**Semantic visual localization.** In the context of long-time operation in dynamic environments, the appearance of a scene can change drastically over time [57, 62, 72]. However, the semantic description of each scene part remains mostly invariant to such changes. Semantic visual localization approaches [5–7, 22, 34, 50, 52, 62, 63, 68, 70, 74, 75, 81, 82] thus use scene understanding, *e.g.*, via semantic segmentation or object detection, as some form of invariant scene representation. Previous work has focused on improving the feature detection and description [34, 62], fea-

<sup>1</sup><http://www.ok.sc.e.titech.ac.jp/res/RIGHTP/>

ture association [7, 37, 52, 63, 70, 75, 81], image retrieval [5, 34, 62, 68, 74, 82], and pose estimation stages [50, 70, 74, 75]. In contrast, this paper focuses on the pose verification stage.

**Pose verification.** Most similar to this paper are works on camera pose verification. The classical approach is to select the pose with the largest number of (weighted) inliers among all candidate poses [25, 29, 61]. However, a (weighted) inlier count is not an appropriate decision criterion in scenes with repetitive structures as an incorrect pose might have more inliers than the correct one [55]. Instead, it is necessary to explicitly account for such structures [55]. Still, focusing on the geometric consistency of feature matches only accounts for positive evidence. In order to take all pixels into account, [72] propose to re-render the scene from the estimated pose. They compare the resulting image with the original query photo using densely extracted RootSIFT [2, 43] features. In this paper, we build on their Dense Pose Verification (DensePV) approach and integrate additional modalities (surface normals and semantic segmentation) into the verification process. These additional modalities further improve the performance of the pose verification stage. While DensePV is a hand-crafted approach, we also propose a trainable variant.

[21, 22, 74, 75] use semantic scene understanding for pose verification: given a pose, they project the 3D points in a scene model into a semantic segmentation of the query image. They measure semantic consistency as the percentage of 3D points projected into an image region with the correct label. Besides determining whether an estimated camera pose is consistent with the scene geometry [21, 22], this measure can be used to identify incorrect matches [75] and to refine pose estimates [70, 74]. We show that this approach, which has so far been used in outdoor scenes, is not applicable in the indoor scenarios considered in this paper.

**View synthesis.** Following [72], we use view synthesis to verify estimated camera poses by re-rendering the scene from the estimated viewpoints. View synthesis has also been used to enable localization under strong appearance [8, 77] or viewpoint [64, 66, 77] changes. However, we are not aware of any previous work that combines multiple modalities and proposes a trainable verification approach.

### 3. Geometric-Semantic Pose Verification

In this paper, we are interested in analyzing the benefits of using more information than pure appearance for camera pose verification in indoor scenes. As such, we propose multiple approaches for pose verification based on the combination of appearance, scene geometry, and semantic information. We integrate our approach into the InLoc pipeline [72], a state-of-the-art visual localization approach for large-scale indoor scenes. In Sec. 3.1, we first review the InLoc algorithm. Sec. 3.2 then discusses how additional

geometric information can be integrated into InLoc’s pose verification stage. Similarly, Sec. 3.3 discusses how semantic information can be used for pose verification.

Since obtaining large training datasets for indoor scenes can be hard, this section focuses on verification algorithms that do not require training data. Sec. 4 then introduces a trainable verification approach.

#### 3.1. Indoor Localization with Pose Verification

The InLoc pipeline represents the scene through a set of RGB-D images with known poses. Given an RGB image as an input query, it first identifies a set of locations in the scene potentially visible in the query via image retrieval. For each location, it performs feature matching and re-ranks the locations based on the number of matches passing a 2D geometric verification stage. Camera poses are then estimated and verified for the top-ranked locations only.

**Candidate location retrieval.** InLoc uses the NetVLAD [1] descriptor to identify the 100 database images most visually similar to the query. For retrieval, the depth maps available for each database image are ignored.

**Dense feature matching and pose estimation (DensePE).** NetVLAD aggregates densely detected CNN features into a compact image-level descriptor. Given the top 100 retrieved images, InLoc performs mutually nearest neighbor matching of the densely extracted CNN features and performs spatial verification by fitting homographies. For the top 10 candidates with the largest number of homography-inliers, InLoc estimates a 6DoF camera pose: The dense 2D-2D matches between the query image and a retrieved database image define a set of 2D-3D matches when taking the depth map of the database image into account. The pose is then estimated using standard P3P-RANSAC [25].

**Dense pose verification (DensePV).** In its final stage, InLoc selects the “best” among the 10 estimated camera poses. To this end, InLoc re-renders the scene from each estimated pose using the color and depth information of the database RGB-D scan: The colored point cloud corresponding to the database RGB-D panoramic scan from which the retrieved database image  $\mathcal{D}$  originated is projected into the estimated pose of the query image  $\mathcal{Q}$  to form a synthetic query image  $\mathcal{Q}_{\mathcal{D}}$ . InLoc’s dense pose verification stage then densely extracts RootSIFT [2, 43] descriptors from both the synthetic and the real query image<sup>2</sup>. It then evaluates the (dis)similarity between the two images as the median of the inverse Euclidean distance between descriptors corresponding to the same pixel position. Let

$$S_D(x, y, \mathcal{D}) = \|\mathbf{d}(\mathcal{Q}, x, y) - \mathbf{d}(\mathcal{Q}_{\mathcal{D}}, x, y)\|^{-1} \quad (1)$$

be the local descriptor similarity function between RootSIFT descriptors extracted at pixel position  $(x, y)$  in  $\mathcal{Q}$  and

<sup>2</sup>RootSIFT is used for robustness to uniform illumination changes.



$\mathcal{Q}_D$ . The similarity score between  $\mathcal{Q}$  and  $\mathcal{Q}_D$  then is

$$\text{DensePV}(\mathcal{Q}, \mathcal{Q}_D) = \text{median}_{x,y}(S_D(x, y, \mathcal{D})) . \quad (2)$$

The median is used instead of the mean as it is more robust to outliers. Invalid pixels, *i.e.*, pixels into which no 3D point projects, are not considered in Eq. 2.

InLoc finally selects the pose estimated using database image  $\mathcal{D}$  that maximizes  $\text{DensePV}(\mathcal{Q}, \mathcal{Q}_D)$ .

### 3.2. Integrating Scene Geometry

Eq. 2 measures the similarity in appearance between the original query image and its synthesized version. The original formulation in the InLoc pipeline has two drawbacks: 1) It only considers the 3D geometry seen from a single scan location corresponding to the retrieved database image  $\mathcal{D}$ . As the pose of the query image can substantially differ from the pose of the database image, this can lead to large regions in the synthesized image into which no 3D points are projected (*c.f.* Fig. 1(i)). 2) Indoor scenes are often dominated by large untextured parts such as white walls. The image appearance of these regions remains constant even under strong viewpoint changes. As such, considering only image appearance in these regions does not provide sufficient information for pose selection. In the following, we propose strategies to address these problems.

**Merging geometry through scan-graphs.** To avoid large regions of missing pixels in the synthesized images, we use 3D data from multiple database RGB-D scans when re-rendering the query. We construct an *image-scan-graph* (*c.f.* Fig. 2) that describes which parts of the scene are related to each database image and are thus used for generating the synthetic query image. Given a retrieved database image  $\mathcal{D}$ , the graph enables us to re-render the query view using more 3D points than those visible in the panoramic RGB-D scan associated with  $\mathcal{D}$ <sup>3</sup>. To construct the graph, we first select the ten spatially closest RGB-D panoramic scans for each database image. We estimate the visibility of each 3D scan in the database image by projecting the 3D points into it while handling occlusions via depth and normal information. We establish a graph edge between the database image and a scan if more than 10% of the database image pixels share the 3D points originating from the scan.

Given a query image  $\mathcal{Q}$ , the retrieved database image  $\mathcal{D}$ , and the estimated camera pose obtained using **DensePE**, we can leverage the constructed scan-graph to render multiple synthetic query images, one for each scan connected to  $\mathcal{D}$  in the graph. These views are then combined by taking depth and normal directions into account to handle occlusions.

Our approach assumes that the scans are dense and rather complete, and that different scans are registered accurately

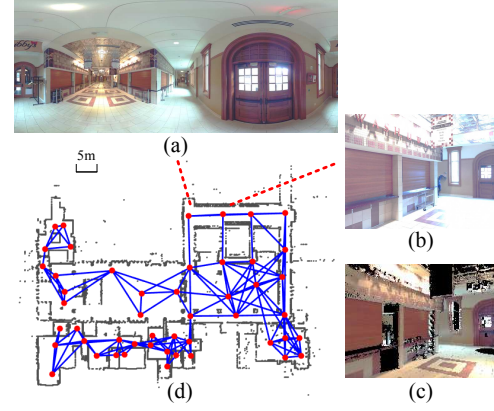


Figure 2: **Image-scan-graph for the InLoc dataset [72].** (a) Example RGB-D panoramic scan. (b) Neighboring database image. (c) 3D points of the RGB-D panoramic scan projected onto the view of the database image. (d) Red dots show where RGB-D panoramic scans are captured. Blue lines indicate links between panoramic scans and database images, established based on visual overlap.

w.r.t. each other. These assumptions do not always hold in practice. Yet, our experiments show that using the scan-graph improves localization performance by reducing the number of invalid pixels in synthesized views compared to using individual scans (*c.f.* Sec. 5).

**Measuring surface normal consistency.** The problem of a lack of information in weakly textured regions can also be addressed by considering other complementary image modalities, such as surface normals. When rendering the synthetic view, we can make use of the depth information in the RGB-D images to create a normal map with respect to a given pose. For each 3D point  $P$  that projects into a 2D point  $p$  in image space, the normal vector is computed by fitting a plane in a local 3D neighbourhood. This 3D neighborhood is defined as the set of 3D points that project within a  $5 \times 5$  pixel patch around  $p$ . This results in a normal map  $\mathcal{N}_D$  as seen from the pose estimated via the database image  $\mathcal{D}$ , where each entry  $\mathcal{N}_D(x, y)$  corresponds to a unit-length surface normal direction. On the query image side, we use a neural network [83] to predict a surface normal map  $\mathcal{N}_Q$ .

We define two verification approaches using surface normal consistency. Both are based on the cosine similarity between normals estimated at pixel position  $(x, y)$ :

$$S_N(x, y, \mathcal{D}) = \mathcal{N}_Q(x, y)^\top \mathcal{N}_D(x, y) . \quad (3)$$

The first strategy, termed *dense normal verification* (**DenseNV**), mirrors DensePV but considers the normal similarities  $S_N$  instead of the descriptor similarities  $S_D$ :

$$\text{DenseNV}(\mathcal{Q}, \mathcal{Q}_D) = \text{median}_{x,y}(S_N(x, y, \mathcal{D})) . \quad (4)$$

The surface normal similarity maps  $S_N$  can contain richer information than the descriptor similarity maps  $S_D$  in the

<sup>3</sup>The InLoc dataset used in our experiments consists of multiple panoramic RGB-D scans, subdivided into multiple database images each.

case of untextured regions. Yet, the contrary will be the case for highly textured regions. Therefore, we propose a second strategy (**DensePNV**), which includes surface normal consistency as a weighting term for the descriptor similarity:

$$\text{DensePNV}(\mathcal{Q}, \mathcal{Q}_D) = \text{median}_{x,y}(w(x, y, \mathcal{D}) \cdot S_D(x, y, \mathcal{D})), \quad (5)$$

where the weighting term  $w(\mathcal{D})$  shifts and normalizes the normal similarities as

$$w(x, y, \mathcal{D}) = \frac{1 + \max(0, S_N(x, y, \mathcal{D}))}{2}. \quad (6)$$

Through  $w$ , the normal similarities act as an attention mechanism on the descriptor similarities, focusing the attention on image regions where normals are consistent.

**Implementation details.** For the query images, for which no depth information is available, surface normals are estimated using [83]. The original implementation from [83] first crops the input image into a square shape and rescales it to  $256 \times 256$  pixels. However, the cropping operation can decrease the field of view and thus remove potentially important information [73]. To preserve the field of view, we modified the network configuration to predict surface normals for rectangular images and scale each image such that its longer side is 256 pixels.

### 3.3. Integrating Scene Semantics

DensePV, DenseNV, and DensePNV implicitly assume that the scene is static, *i.e.*, that the synthesized query image should look identical to the real query photo. In practice, this assumption is often violated as scenes change over time. For example, posters on walls or bulletin boards might be changed or furniture might be moved around. Handling such changes requires a higher-level understanding of the scene, which we model via semantic scene understanding.

**Projective Semantic Consistency (PSC).** A standard approach to using scene understanding for pose verification is to measure semantic consistency [21, 74, 75]: These methods use a semantically labeled 3D point cloud, *e.g.*, obtained by projecting semantic labels extracted from RGB images onto a point cloud, and a semantic segmentation of the query image. The labeled 3D point cloud is projected into the query image via an estimated pose. Semantic consistency is then computed by counting the number of matching labels between the query and the synthetic image.

**Ignoring transient objects.** PSC works well in outdoor scenes, where there are relatively many classes and where points projecting into “empty” regions such as sky clearly indicate incorrect / inaccurate pose estimates. Yet, we will show in Sec. 5 that it does not work well in indoor scenes. This is due to the fact that there are no “empty” regions and that most pixels belong to walls, floors, or ceilings. Instead

of enforcing semantic consistency everywhere, we use semantic information to determine where we expect geometric and appearance information to be unreliable.

We group the semantic classes into five “superclasses”: *people*, *transient*, *stable*, *fixed*, and *outdoor*. The *transient* superclass includes easily-movable objects, *e.g.*, chairs, books, or trash cans. The *stable* superclass contains objects that are moved infrequently, *e.g.*, tables, couches, or wardrobes. The *fixed* superclass contains objects that are unlikely to move, *e.g.*, walls, floors, and ceilings. When computing DensePV, DenseNV, or DensePNV scores, we ignore pixels in the query image belonging to the *people* and *transient* superclasses. We refer to these approaches as **DensePV+S**, **DenseNV+S**, and **DensePNV+S**.

**Implementation details.** Semantics are extracted using the CSAIL Semantic Segmentation/Scene Parsing approach [87, 88] based on a Pyramid Scene Parsing Network [86], trained on the ADE20K dataset [87, 88] containing 150 classes. Details on the mapping of classes to superclasses are provided in [73].

## 4. Trainable Pose Verification

In the previous section, we developed several methods for camera pose verification that did not require any training data. Motivated by the recent success of trainable methods for several computer vision tasks, this section presents a trainable approach for pose verification (**TrainPV**), where we will train a pose verification scoring function from examples of correct and incorrect poses. We first describe the proposed model (*c.f.* Fig. 3), then how we obtained training data, and finally the loss used for training.

**Network architecture for pose verification.** Our network design follows an approach similar to that of DensePV, where given the original  $\mathcal{Q}$  and a synthetic query image  $\mathcal{Q}_D$  we first extract dense feature descriptors  $\mathbf{d}(\mathcal{Q}, x, y)$  and  $\mathbf{d}(\mathcal{Q}_D, x, y)$  using a fully convolutional network. This feature extraction network plays the role of the dense RootSIFT descriptor of DensePV. Then, a descriptor similarity score map is computed by the cosine similarity<sup>4</sup>:

$$S_D(x, y, \mathcal{D}) = \mathbf{d}(\mathcal{Q}, x, y)^\top \mathbf{d}(\mathcal{Q}_D, x, y). \quad (7)$$

Finally, the 2D descriptor similarity score-map given by Eq. 7 is processed by a *score regression CNN* that estimates the agreement between  $\mathcal{Q}$  and  $\mathcal{Q}_D$ , resulting in a scalar score. This score regression CNN is composed of several convolution layers followed by ReLU non-linearities and a final average pooling layer. The intuition is that the successive convolution layers can identify coherent similarity (and dissimilarity) patterns in the descriptor similarity score-map

<sup>4</sup>The descriptors are L2 normalized beforehand.




Figure 3: **Network architecture for Trainable Pose Verification.** Input images are passed through a feature extraction network  $F$  to obtain dense descriptors  $\mathbf{f}$ . These are then combined by computing the descriptor similarity map  $S_D$ . Finally a score regression CNN  $R$  produces the score  $s$  of the trainable pose verification model.

$S_D$ . A final average pooling then aggregates the positive and negative evidence over the score map to accept or reject the candidate pose. Note that our architecture bears a resemblance to recent methods for image matching [51] and optical flow estimation [24]. Contrary to these methods, which estimate global geometric transformations or local displacement fields, our input images ( $Q$ ,  $Q_D$ ) are already spatially aligned and we seek to measure their agreement.

**Training data.** In order to train the proposed network, we need appropriate annotated training data. For this, we use additional video sequences recorded for the InLoc benchmark [72], which are separate from the actual test images. We created 6DoF camera poses for the new images via manual annotation and Structure-from-Motion (SfM) [73]. For each image, we generated pose candidates for training in two different ways.

The first approach randomly perturbs the ground-truth pose with 3D translations and rotations up to  $\pm 1\text{m}$  and  $\pm 20\text{deg}$ . We use the perturbed random poses to generate synthetic images by projecting the 3D point cloud from the InLoc database scan associated to that image.

The second approach uses the DensePE pipeline [72] as a way of generating realistic estimated poses for the additional images. For this, we run the images through the localization pipeline, obtaining pose estimates and the corresponding database images. Then we run synthetic image rendering on these poses and use these images for training. Note that, contrary to the randomized approach where images are generated from the correct query-database image pair, here the synthetic images might be generated from unrelated pairs. This is because the localization pipeline might fail to generate “hard-negatives”, *i.e.*, examples corresponding to other similar-looking but different locations.

In both cases, for each real and synthetic image pair, both the ground-truth ( $P_{GT}$ ) and the estimated ( $\tilde{P}$ ) poses are known. In order to generate a scalar score that can be used as a training signal, we compute the mean 2D reprojection error  $r(Q_D)$  of the 3D point cloud  $\{X_i\}_{i=1}^N$  in image space:

$$r(Q_D) = \frac{1}{N} \sum_{i=1}^N \|\mathcal{P}(X_i, P_{GT}) - \mathcal{P}(X_i, \tilde{P})\|, \quad (8)$$

where  $\mathcal{P}$  is the 3D-2D projection function.

**Training loss.** A suitable loss is needed in order to train the above network for pose verification. Given a query image  $Q$  and a set of candidate synthesized query images  $\{Q_{D_i}\}_{i=1}^N$ , we would like to re-rank the candidates in the order given by the average reprojection errors (*c.f.* Eq. 8).

In order to do this, we assume that each synthetic image  $Q_{D_i}$  has an associated discrete probability  $p(Q_{D_i})$  of corresponding to the best matching pose for the query image  $Q$  among the  $N$  candidates. This probability should be inversely related to the reprojection error from Eq. 8, *i.e.*, a pose with a high reprojection error has little probability of being the best match. Then, the scores  $s_i = s(Q, Q_{D_i})$  produced by our trained pose verification CNN can be used to model an estimate of this probability  $\hat{p}(Q_{D_i})$  as

$$\hat{p}(Q_{D_i}) = \frac{\exp(s_i)}{\sum_{k=1}^N \exp(s_k)}. \quad (9)$$

To define the ground-truth probability distribution  $p$ , we make use of the reprojection error  $r_i = r(Q_{D_i})$  from Eq. 8:

$$p(Q_{D_i}) = \frac{\exp(-\tilde{r}_i)}{\sum_{k=1}^N \exp(-\tilde{r}_k)}, \quad (10)$$

where  $\tilde{r}_i = r_i / \min_k r_k$  is the relative reprojection error with respect to the minimum value within the considered candidates. The soft-max function is used to obtain a normalized probability distribution<sup>5</sup>.

The training loss  $\mathcal{L}$  is defined as the cross-entropy between the ground-truth and estimated distributions  $p$  and  $\hat{p}$ :

$$\mathcal{L} = - \sum_{i=1}^N p(Q_{D_i}) \log \hat{p}(Q_{D_i}), \quad (11)$$

where the sum is over the  $N$  candidate poses.

Note that because the ground-truth score distribution  $p$  is fixed, minimizing the cross-entropy between  $p$  and  $\hat{p}$  is equivalent to minimizing the Kullback-Leibler divergence between these two distributions. Thus, the minimum is achieved when  $\hat{p}$  matches  $p$  exactly. Also note that, at the optimum, the ground-truth ranking between the candidate poses is respected, as desired.

**Implementation details.** The feature extraction network is composed of a fully convolutional ResNet-18 architecture (up to the conv4-2 layer) [28], pretrained on ImageNet [23]. Its weights are kept fixed during training as the large number of parameters would lead to overfitting in our small-sized training sets. The score regression CNN

<sup>5</sup>Relative reprojection errors are used to prevent saturation of the soft-max function.

is composed of four convolutional layers with  $5 \times 5$  filters and a padding of 2, each followed by ReLU non-linearities. Each convolutional layer operates on 32 channels as input and output, except the first one, which takes the single channel descriptor similarity map  $d_D$  as input, and the last one, which also outputs a single channel tensor. Finally, an average pooling layer is used to obtain the final score estimate  $s(Q, Q_D)$ . The score regression CNN is trained for 10 epochs using the PyTorch framework [48], with the Adam optimizer and a learning rate of  $10^{-5}$ .

## 5. Experimental Evaluation

**Dataset.** We evaluate our approach in the context of indoor visual localization on the recently proposed InLoc dataset [72]. The dataset is based on the 3D laser scan model from [80] and depicts multiple floors in multiple university buildings. The 10k database images correspond to a set of perspective images created from RGB-D panoramic scans captured using a camera mounted on a laser scanner, *i.e.*, a depth map is available for each database image. The 329 query images were recorded using an iPhone7 about a year after the database images and at different times of the day compared to the database images. The resulting changes in scene appearance between the query and database images make the dataset significantly more challenging than other indoor datasets such as 7 Scenes [65].

**Evaluation measure.** Following [57, 72], we measure the errors of an estimated pose as the differences in position and orientation from the reference pose provided by the dataset. We report the percentage of query images whose estimated poses differ by no more than  $X$  meters and  $Y$  degrees from the reference pose for different pairs of thresholds  $(X, Y)$ .

**Baselines.** Our approach is based on the InLoc approach [72], which is the current state-of-the-art in large-scale indoor localization and thus serves as our main baseline (*c.f.* Sec. 3.1). We build on top of the code released by the authors of [72]. For a given input image, we first retrieved the top 100 database images via NetVLAD with a pre-trained Pitts30K [1] VGG-16 [67] model. Feature matching is then performed between query and retrieved images also using the densely extracted CNN features of NetVLAD’s VGG-16 [67] architecture. After re-ranking the image list according to the number of homography-inliers, we estimate pose candidates for the top 10 best matched images using a set of dense inlier matches and database depth information (DensePE).

For each candidate, **DensePV** renders the view with respect to the RGB-D panoramic scan from which the database image originated. The similarity between the original and the rendered view is computed as the median distance of densely extracted hand-crafted features [2, 43].

As a semantic baseline, we project the database 3D

Method	Error [meter, degree]			
	[0.25, 5]	[0.50, 5]	[1.00, 10]	[2.00, 10]
<b>w/o scan-graph</b>				
DensePE [72]	35.0	46.2	57.1	61.1
DensePV [72]	38.9	55.6	69.9	74.2
PSC	30.4	44.4	55.9	58.4
DensePV+S	39.8	57.8	71.1	75.1
DenseNV	32.2	45.6	58.1	62.9
DenseNV+S	31.6	46.5	60.5	64.4
DensePNV	40.1	58.1	72.3	<b>76.6</b>
DensePNV+S	40.1	59.0	72.6	76.3
<b>w/ scan-graph</b>				
DensePV	39.8	59.0	69.0	71.4
PSC	28.3	43.2	55.0	58.4
DensePV+S	<b>41.3</b>	<b>61.7</b>	71.4	74.2
DenseNV	34.3	50.5	62.9	66.6
DenseNV+S	35.9	51.4	64.4	68.4
DensePNV	40.4	60.5	<b>72.9</b>	75.4
DensePNV+S	41.0	60.5	72.3	75.1
TrainPV (random)	39.5	56.5	72.3	76.3
TrainPV (DPE)	39.5	56.8	72.3	76.3
Oracle (Upper-bound)	43.5	63.8	77.5	80.5

Table 1: **The impact of using the scan-graph for pose verification evaluated on the InLoc dataset [72].** We report the percentage of queries localized within given positional and rotational error bounds.

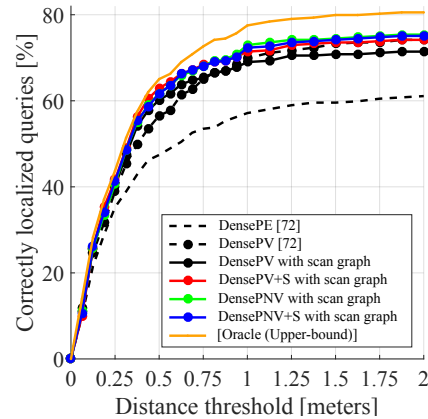


Figure 4: **The impact of geometric and semantic information on the pose verification stage.** We validate the performance of the proposed methods that consider additional geometric and semantic information on the InLoc dataset [72]. Each curve shows the percentage of the queries localized within varying distance thresholds (x-axis) and a fixed rotational error of at most 10 degrees.

points, labeled via the database image, into the query and count the number of points with consistent labels (**PSC**).

As reported in [72], both DSAC [10, 11] and PoseNet [32, 33] fail to train on the InLoc dataset. We thus do not consider them in our experiments.

**The impact of using additional modalities.** Tab. 1 and Fig. 4 compare the localization performance of the baseline pose verification methods against our novel variants proposed in Sec. 3. DenseNV and PSC perform worst, even compared to the baseline DensePE. This is not surprising as both completely ignore the visual appearance and instead focus on information that by itself is less discriminative






Figure 5: **Typical failure cases of view synthesis using the scan-graph.** Top: Synthetic images obtained during DensePV with the scan-graph, affected by (a) misalignment of the 3D scans to the floor plan, (b) sparsity of the 3D scans, and (c) intensity changes. Bottom: A typical failure case of DensePV with the scan-graph: (d) query image, (e) re-rendered query, (f) error map computed with RootSIFT.

(surface normals and semantic segmentation, respectively). On the other hand, combining geometric and / or semantic information with appearance information improves the localization performance compared to DensePV. This clearly validates our idea of using multiple modalities.

We observe the biggest improvement by using our scan-graph, which is not surprising as it reduces the number of invalid pixels and thus adds more information to the rendered images. DensePV+S using a scan-graph shows the best performance at higher accuracy levels. DensePNV using the scan-graph combines appearance and normal information and constantly shows more than a 5% performance gain compared to DensePV. Yet, DensePNV+S with the scan-graph shows less improvement compared to its single scan variant and even performs worse for larger error thresholds. This is partially due to inaccurate depths and camera poses of the database images (*c.f.* Fig. 5 (a–c)). There are also failures where a single scan already provides a rather complete view. Fig. 5 (d–f) shows such an example: due to weak textures, the rendering appears similar to the query image. Such failures cannot be resolved using the scan-graph.

Interestingly, simply combining all modalities does not necessarily lead to the best performance. To determine whether the modalities are simply not complementary or whether this is due to the way they are combined, we create an oracle. The oracle is computed from four of our proposed variants (DensePV [72], DensePV w/ scan-graph, DensePV+S w/ scan graph, and DensePNV w/ scan-graph): Each variant provides a top-ranked pose and the oracle, having access to the ground truth, simply selects the pose with the smallest error. As can be seen from Tab. 1 and Fig. 4, the oracle performs clearly better than any of our proposed variants. We also observed DenseNV+S provides better poses than the oracle (which does not use DenseNV+S) for

about 9% the queries, which could lead to further improvements. This shows that the different modalities are indeed complementary. Therefore, we attribute the diminishing returns observed for DensePNV+S to the way we combine semantic and normal information. We assume that better results could be obtained with normals *and* semantics if one reasons about the consistency of image regions rather than on a pixel level (as is done by using the median).

**Trainable pose verification.** We next evaluate two trainable approaches (TrainPV), which are trained by randomly perturbed views (random) or by selecting views based on DensePE estimation (DPE) (*c.f.* Sec. 4). Even though both are trained using only appearance information, they still are able to use higher-level scene context as they use dense features extracted from a pre-trained fully convolutional network. Tab. 1 compares both TrainPV variants with the baselines and our hand-crafted approaches. Even though both variants use different training sets, they achieve nearly the same performance<sup>6</sup>. This indicates that the choice of training set is not critical in our setting. The results show that TrainPV outperforms the DensePV baseline, but not necessarily our hand-crafted variants based on multiple modalities. This result validates our idea of pose verification based on different sources of information. We also tried variants of TrainPV that use multiple modalities, but did not observe further improvements.

## 6. Conclusion

We have presented a new pose verification approach to improve large-scale indoor camera localization, which is extremely challenging due to the existence of repetitive structures, weakly-textured scenes, and dynamically appearing/disappearing objects over time. To address these challenges, we have developed and validated multiple strategies to combine appearance, geometry, and semantics for pose verification, showing significant improvements over a current state-of-the-art indoor localization baseline. To encourage further progress on the challenging indoor localization problem, we make our code publicly available.

**Acknowledgments.** This work was partially supported by JSPS KAKENHI Grant Numbers 15H05313, 17H00744, 17J05908, EU-H2020 project LADIO No. 731970, ERC grant LEAP No. 336845, CIFAR Learning in Machines & Brains program and the EU Structural and Investment Funds, Operational Programme Research, Development and Education under the project IMPACT (reg. no. CZ.02.1.01/0.0/0.0/15\_003/0000468). We gratefully acknowledge the support of NVIDIA Corporation with the donation of Quadro P6000 GPU.

<sup>6</sup>We are considering a discrete re-ranking problem on a few candidate poses per query. As such, it is not surprising to have very similar results.



## References

- [1] Relja Arandjelović, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *Proc. CVPR*, 2016. 2, 3, 7
- [2] Relja Arandjelović and Andrew Zisserman. Three things everyone should know to improve object retrieval. In *Proc. CVPR*, 2012. 3, 7
- [3] Relja Arandjelovic and Andrew Zisserman. All about VLAD. In *Proc. CVPR*, 2013. 2
- [4] Relja Arandjelović and Andrew Zisserman. Dislocation: Scalable descriptor distinctiveness for location recognition. In *Proc. ACCV*, 2014. 2
- [5] Relja Arandjelović and Andrew Zisserman. Visual vocabulary with a semantic twist. In *Proc. ACCV*, 2014. 2, 3
- [6] Shervin Ardeshtir, Amir Roshan Zamir, Alejandro Torroella, and Mubarak Shah. GIS-Assisted Object Detection and Geospatial Localization. In *Proc. ECCV*, 2014. 2
- [7] Nikolay Atanasov, Menglong Zhu, Kostas Daniilidis, and George J. Pappas. Localization from semantic observations via the matrix permanent. *Intl. J. of Robotics Research*, 35(1-3):73–99, 2016. 2, 3
- [8] Mathieu Aubry, Bryan C. Russell, and Josef Sivic. Painting-to-3D Model Alignment via Discriminative Visual Elements. *ACM Trans. Graph.*, 33(2):14:1–14:14, Apr 2014. 3
- [9] Vassileios Balntas, Shuda Li, and Victor Adrian Prisacariu. RelocNet: Continuous Metric Learning Relocalisation using Neural Nets. In *Proc. ECCV*, 2018. 2
- [10] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. DSAC - Differentiable RANSAC for Camera Localization. In *Proc. CVPR*, 2017. 1, 2, 7
- [11] Eric Brachmann and Carsten Rother. Learning less is more-6D camera localization via 3d surface regression. In *Proc. CVPR*, 2018. 1, 2, 7
- [12] Samarth Brahmabhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-Aware Learning of Maps for Camera Localization. In *Proc. CVPR*, 2018. 2
- [13] Song Cao and Noah Snavely. Graph-based discriminative learning for location recognition. In *Proc. CVPR*, 2013. 2
- [14] Song Cao and Noah Snavely. Minimal scene descriptions from structure from motion models. In *Proc. CVPR*, 2014. 2
- [15] Robert Castle, Georg Klein, and David W. Murray. Video-rate localization in multiple maps for wearable augmented reality. In *ISWC*, 2008. 1
- [16] Tommaso Cavallari, Stuart Golodetz, Nicholas A. Lord, Julien Valentin, Luigi Di Stefano, and Philip H. S. Torr. On-The-Fly Adaptation of Regression Forests for Online Camera Relocalisation. In *Proc. CVPR*, 2017. 1, 2
- [17] David M. Chen, Georges Baatz, Kevin Köser, Sam S Tsai, Ramakrishna Vedantham, Timo Pylvänäinen, Kimmo Roimela, Xin Chen, Jeff Bach, Marc Pollefeys, et al. City-scale landmark identification on mobile devices. In *Proc. CVPR*, 2011. 2
- [18] Ondřej Chum and Jiří Matas. Matching with PROSAC-progressive sample consensus. In *Proc. CVPR*, 2005. 2
- [19] Ondřej Chum and Jiří Matas. Optimal randomized RANSAC. *IEEE PAMI*, 30(8):1472–1482, 2008. 2
- [20] Ondřej Chum, Andrej Mikulík, Michal Perdoch, and Jiří Matas. Total recall II: Query expansion revisited. In *Proc. CVPR*, 2011. 2
- [21] Andrea Cohen, Torsten Sattler, and Mark Pollefeys. Merging the Unmatchable: Stitching Visually Disconnected SfM Models. In *Proc. ICCV*, 2015. 2, 3, 5
- [22] Andrea Cohen, Johannes Lutz Schönberger, Pablo Speciale, Torsten Sattler, Jan-Michael Frahm, and Marc Pollefeys. Indoor-Outdoor 3D Reconstruction Alignment. In *Proc. ECCV*, 2016. 2, 3
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. CVPR*, 2009. 6
- [24] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proc. ICCV*, pages 2758–2766, 2015. 6
- [25] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Comm. ACM*, 24(6):381–395, 1981. 2, 3
- [26] Petr Gronat, Guillaume Obozinski, Josef Sivic, and Tomas Pajdla. Learning and calibrating per-location classifiers for visual place recognition. In *Proc. CVPR*, 2013. 2
- [27] Bert M. Haralick, Chung-Nan Lee, Karsten Ottenberg, and Michael Nölle. Review and analysis of solutions of the three point perspective pose estimation problem. *IJCV*, 13(3):331–356, 1994. 2
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, 2016. 6
- [29] Arnold Irschara, Christopher Zach, Jan-Michael Frahm, and Horst Bischof. From Structure-from-Motion point clouds to fast location recognition. In *Proc. CVPR*, 2009. 2, 3
- [30] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *Proc. ECCV*, 2008. 2
- [31] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Packing bag-of-features. In *Proc. ICCV*, 2009. 2
- [32] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proc. CVPR*, 2017. 2, 7
- [33] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *Proc. ICCV*, 2015. 2, 7
- [34] Hyo Jin Kim, Enrique Dunn, and Jan-Michael Frahm. Learned contextual feature reweighting for image geolocalization. In *Proc. CVPR*, 2017. 2, 3
- [35] L. Kneip, D. Scaramuzza, and R. Siegwart. A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In *Proc. CVPR*, 2011. 2
- [36] Jan Knopp, Josef Sivic, and Tomas Pajdla. Avoiding confusing features in place recognition. In *Proc. ECCV*, 2010. 2

- [37] Nikolay Kobyshev, Hayko Riemenschneider, and Luc Van Gool. Matching Features Correctly through Semantic Understanding. In *Proc. 3DV*, 2014. 3
- [38] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. Real-Time Solution to the Absolute Pose Problem with Unknown Radial Distortion and Focal Length. In *Proc. ICCV*, 2013. 2
- [39] Yunpeng Li, Noah Snavely, and Daniel P. Huttenlocher. Location recognition using prioritized feature matching. In *Proc. ECCV*, 2010. 2
- [40] Yunpeng Li, Noah Snavely, Daniel P. Huttenlocher, and Pascal Fua. Worldwide pose estimation using 3d point clouds. In *Proc. ECCV*, 2012. 1, 2
- [41] Hyon Lim, Sudipta N. Sinha, Michael F. Cohen, and Matthew Uyttendaele. Real-time image-based 6-DOF localization in large-scale environments. In *Proc. CVPR*, 2012. 1
- [42] Liu Liu, Hongdong Li, and Yuchao Dai. Efficient Global 2D-3D Matching for Camera Localization in a Large-Scale 3D Map. In *Proc. ICCV*, 2017. 2
- [43] David G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004. 2, 3, 7
- [44] Simon Lynen, Torsten Sattler, Michael Bosse, Joel A. Hesch, Marc Pollefeys, and Roland Siegwart. Get Out of My Lab: Large-scale, Real-Time Visual-Inertial Localization. In *Proc. RSS*, 2015. 1
- [45] Daniela Massiceti, Alexander Krull, Eric Brachmann, Carsten Rother, and Philip H.S. Torr. Random Forests versus Neural Networks - What's Best for Camera Relocalization? In *Proc. Intl. Conf. on Robotics and Automation*, 2017. 2
- [46] Lili Meng, Jianhui Chen, Frederick Tung, James J. Little, Julien Valentin, and Clarence W. de Silva. Backtracking Regression Forests for Accurate Camera Relocalization. In *Proc. IEEE/RSJ Conf. on Intelligent Robots and Systems*, 2017. 1, 2
- [47] Lili Meng, Frederick Tung, James J. Little, Julien Valentin, and Clarence W. de Silva. Exploiting Points and Lines in Regression Forests for RGB-D Camera Relocalization. In *Proc. IEEE/RSJ Conf. on Intelligent Robots and Systems*, 2018. 1, 2
- [48] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017. 7
- [49] James Philbin, Ondřej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *Proc. CVPR*, 2007. 1, 2
- [50] Noha Radwan, Abhinav Valada, and Wolfram Burgard. VLocNet++: Deep multitask learning for semantic visual localization and odometry. *IEEE Robotics And Automation Letters (RA-L)*, 3(4):4407–4414, 2018. 2, 3
- [51] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *Proc. CVPR*, 2017. 6
- [52] Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H. J. Kelly, and Andrew J. Davison. SLAM++: Simultaneous Localisation and Mapping at the Level of Objects. In *Proc. CVPR*, 2013. 2, 3
- [53] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proc. CVPR*, 2019. 2
- [54] Torsten Sattler, Michal Havlena, Filip Radenovic, Konrad Schindler, and Marc Pollefeys. Hyperpoints and fine vocabularies for large-scale location recognition. In *Proc. ICCV*, 2015. 2
- [55] Torsten Sattler, Michal Havlena, Konrad Schindler, and Marc Pollefeys. Large-scale location recognition and the geometric burstiness problem. In *Proc. CVPR*, 2016. 2, 3
- [56] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Efficient & effective prioritized matching for large-scale image-based localization. *IEEE PAMI*, 39(9):1744–1756, 2017. 1, 2
- [57] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DOF outdoor visual localization in changing conditions. In *Proc. CVPR*, 2018. 2, 7
- [58] Torsten Sattler, Akihiko Torii, Josef Sivic, Marc Pollefeys, Hajime Taira, Masatoshi Okutomi, and Tomas Pajdla. Are large-scale 3D models really necessary for accurate visual localization? In *Proc. CVPR*, 2017. 2
- [59] Torsten Sattler, Qunjie Zhou, Mark Pollefeys, and Laura Leal-Taixé. Understanding the Limitations of CNN-based Absolute Camera Pose Regression. In *Proc. CVPR*, 2019. 2
- [60] Grant Schindler, Matthew Brown, and Richard Szeliski. City-Scale Location Recognition. In *Proc. CVPR*, 2007. 2
- [61] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-From-Motion Revisited. In *Proc. CVPR*, 2016. 1, 2, 3
- [62] Johannes Lutz Schönberger, Marc Pollefeys, Andreas Geiger, and Torsten Sattler. Semantic Visual Localization. In *Proc. CVPR*, 2018. 2, 3
- [63] Markus Schreiber, Carsten Knöppel, and Uwe Franke. Lane-Loc: Lane marking based localization using highly accurate maps. In *Proc. IV*, 2013. 2, 3
- [64] Qi Shan, Changchang Wu, Brian Curless, Yasutaka Furukawa, Carlos Hernandez, and Steven M. Seitz. Accurate Geo-Registration by Ground-to-Aerial Image Matching. In *Proc. 3DV*, 2014. 3
- [65] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in RGB-D images. In *Proc. CVPR*, 2013. 1, 2, 7
- [66] Dominik Sibbing, Torsten Sattler, Bastian Leibe, and Leif Kobbelt. SIFT-realistic rendering. In *Proc. 3DV*, 2013. 3
- [67] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proc. ICLR*, 2015. 7
- [68] Gautam Singh and Jana Koščeká. Semantically Guided Geolocation and Modeling in Urban Environments. In *Large-Scale Visual Geo-Localization*, 2016. 2, 3
- [69] Josef Sivic and Andrew Zisserman. Video google: A text retrieval approach to object matching in videos. In *Proc. ICCV*, 2003. 1, 2
- [70] E. Stenborg, C. Toft, and L. Hammarstrand. Long-term Visual Localization using Semantically Segmented Images. In *Proc. Intl. Conf. on Robotics and Automation*, 2018. 2, 3

- [71] Linus Svärm, Olof Enqvist, Fredrik Kahl, and Magnus Oskarsson. City-Scale Localization for Cameras with Known Vertical Direction. *IEEE PAMI*, 39(7):1455–1461, 2017. 1, 2
- [72] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. InLoc: Indoor visual localization with dense matching and view synthesis. In *Proc. CVPR*, 2018. 1, 2, 3, 4, 6, 7, 8
- [73] Hajime Taira, Ignacio Rocco, Jiri Sedlar, Masatoshi Okutomi, Josef Sivic, Tomas Pajdla, Torsten Sattler, and Akihiko Torii. Is This The Right Place? Geometric-Semantic Pose Verification for Indoor Visual Localization. *arXiv preprint arXiv:1908.04598*, 2019. 5, 6
- [74] Carl Toft, Carl Olsson, and Fredrik Kahl. Long-term 3D Localization and Pose from Semantic Labellings. In *Proc. ICCV Workshops*, 2017. 2, 3, 5
- [75] Carl Toft, Erik Stenborg, Lars Hammarstrand, Lucas Brynte, Marc Pollefeys, Torsten Sattler, and Fredrik Kahl. Semantic Match Consistency for Long-Term Visual Localization. In *Proc. ECCV*, 2018. 2, 3, 5
- [76] Giorgos Tolias and Hervé Jégou. Local visual query expansion: Exploiting an image collection to refine local descriptors. Technical Report RR-8325, INRIA, 2013. 2
- [77] Akihiko Torii, Relja Arandjelovic, Josef Sivic, Masatoshi Okutomi, and Tomas Pajdla. 24/7 place recognition by view synthesis. In *Proc. CVPR*, 2015. 2, 3
- [78] Akihiko Torii, Josef Sivic, Tomas Pajdla, and Masatoshi Okutomi. Visual place recognition with repetitive structures. In *Proc. CVPR*, 2013. 2
- [79] Florian Walch, Caner Hazirbas, Laura Leal-Taixé, Torsten Sattler, Sebastian Hilsenbeck, and Daniel Cremers. Image-Based Localization Using LSTMs for Structured Feature Correlation. In *Proc. ICCV*, 2017. 2
- [80] Erik Wijmans and Yasutaka Furukawa. Exploiting 2D floor-plan for building-scale panorama RGBD alignment. In *Proc. CVPR*, 2017. 7
- [81] Fisher Yu, Jianxiong Xiao, and Thomas A. Funkhouser. Semantic alignment of LiDAR data at city scale. In *Proc. CVPR*, 2015. 2, 3
- [82] X. Yu, S. Chaturvedi, C. Feng, Y. Taguchi, T.-Y. Lee, C. Fernandes, and S. Ramalingam. VLASE: Vehicle Localization by Aggregating Semantic Edges. In *Proc. IEEE/RSJ Conf. on Intelligent Robots and Systems*, 2018. 2, 3
- [83] Amir Roshan Zamir, Alexander Sax, , William B. Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proc. CVPR*, 2018. 4, 5
- [84] Amir Roshan Zamir and Mubarak Shah. Accurate image localization based on google maps street view. In *Proc. ECCV*, 2010. 2
- [85] Bernhard Zeisl, Torsten Sattler, and Marc Pollefeys. Camera Pose Voting for Large-Scale Image-Based Localization. In *Proc. ICCV*, 2015. 1, 2
- [86] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proc. CVPR*, 2017. 5
- [87] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ADE20K dataset. In *Proc. CVPR*, 2017. 5
- [88] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ADE20K dataset. *IJCV*, 127(3):302–321, 2019. 5